Memorandum of Understanding between observatories participating in the Astrophysical Multimessenger Observatory Network

AMON Executive Board

May 24, 2019

The Astrophysical Multimessenger Observatory Network (AMON) provides a framework for correlating high energy astrophysical signals across all possible astronomical messengers: photons, neutrinos, cosmic rays, and gravitational waves. The primary goals of the program are: (1) To allow participating observatories to share their data with one another with strict anonymity, confidentiality and in accordance with their blind analysis procedures, (2) To enhance the combined sensitivity of participating observatories to astrophysical transients by enabling them to search for coincidences in their sub-threshold archival data and then in their sub-threshold real-time data and (3) To enable follow-up imaging of possible astrophysical sources with minimal latency.

Membership

Participants in AMON can be characterized as either "triggering," "follow-up" or both. Triggering participants are generally wide field-of-view observatories that feed a stream of sub-threshold events into the AMON framework. These events are processed to search for temporal and spatial correlations, leading to AMON-generated "alerts." Follow-up participants generally search for electromagnetic counterparts to the alerts with high-throughput, smaller field-of-view instruments. The list of participating observatories is provided on the AMON wiki [1]. All signatories to this MoU are participants in the AMON Consortium.

New triggering or follow-up observatories will be welcome to join AMON. New triggering observatories must undertake a set of studies with simulated or archival data and a trial period with real-time data, demonstrating a potential positive impact on coincidence searches. Similarly, new follow-up observatories will be expected to demonstrate a robust response to simulated alerts.

The corresponding documentation will be circulated among current AMON triggering observatories for a period of two weeks. If, in this period, objections are raised by the current participants, the AMON development team will work with the new observatory to reconcile the issues as quickly as possible.

Operational Phases

There are three distinct operational phases for coincidence searches performed in the AMON framework, with this MoU covering all of them.

1. In the initial development phase there will be a shake-down period during which only simulated or existing archival data will be used. During this time, the participating observatories will establish and test analysis algorithms. Archival analyses follow the same data control and confidentiality procedures as described herein for real-time analyses.

2. The next phase is a real-time phase, when alerts will be mainly distributed in-network (*i.e.*, only to observatories that have signed this MoU). During this period, observatories will be incorporated into the real-time system.

3. In the final phase, it is anticipated that each coincidence search will then enter an open real-time phase, in which alerts are distributed publicly, although in a controlled manner as described later in this MoU.

Confidentiality

All data sent to AMON remains the intellectual property of the contributing observatory and can not be released without its approval. This applies to data released following the guidelines established in this MoU and to data released through any other mechanism internal to AMON participants. As such, all participants agree to keep confidential any and all information they may receive via AMON about another observatory's data.

An AMON-generated alert consists of an ID number, as well as the estimated coincidence time, direction, and positional error ellipse or equivalent, the number of triggering observatories, and a combined measure of significance for the alert. To maintain as much anonymity as possible, initial versions of each alert may not reference the specific triggering observatories, with that information only being released upon approval of the contributing observatories. In addition:

- A non-detection by an observatory that is actively sensitive in the time/position window of the AMON alert is of scientific merit and hence must also be held confidential.

- While the list of triggering observatories is not explicitly included in the initial version of an alert, in some cases it may be inferred from from certain characteristics thereof (*e.g.*, from the error ellipse). Any information that can be inferred from non-public data in such a manner is considered confidential.

- While private communication between participants may include a discussion of event origin, any AMON information that is not in the public domain is considered confidential.

Additional Requirements

Each member collaboration must provide the name of one or several points of contact. Points of contact are responsible for (1) Developing and maintaining this MoU, (2) Representing and reporting back to his/her collaboration on matters related to AMON, and (3) Forming and guiding analysis algorithm working groups from the participating collaborations.

The AMON real-time system requires pre-defined analysis algorithms, are developed by the participating observatories in close consultation with one another. The content of these algorithms is listed on the AMON wiki. Each triggering observatory may appoint individuals to contribute to algorithm development and will be asked to review/approve the analysis algorithms before we enter each major phase of AMON.

The algorithms require specific inputs from each observatory, including pointing information and orientation, event location and position uncertainty (or equivalent), and background event rate. Some of this information will be transmitted to AMON in the form of event summaries (*e.g.*, in VOEvent format [2]), via client software installed on the observatory's computers. Other information may be transmitted periodically via other means and stored as a configuration file in the AMON system. To preserve confidentiality and data analysis blindness within each collaboration, the default is for access to each observatory's data to be limited to a small number of individuals in that collaboration. These individuals will be appointed by their respective collaborations. Each participating observatory agrees to make a best effort to establish efficient communication of these data. Additionally, participants may also be asked to provide and process test events and alerts (events and alerts that are highly unlikely to be of astrophysical origin), time- or position-scrambled archival data, system heartbeats and handshakes.

For each AMON alert issued, the triggering observatories are required to check their data for errors and sign off on it within a specified time period. The follow-up observatories are required to issue a short report (*e.g.*, as a VOEvent) to the AMON network within a predetermined time period after each alert they process, even in the case of a non-detection. This time period may vary as a function of the analysis algorithm underlying the alert. The AMON team will keep track of these time periods on the AMON wiki, accessible to all participating observatories.

Source Discovery

In phases 2 and 3 of AMON, when an alert is determined to be of high significance it will initially be sent only to the participating triggering and follow-up observatories via private channels in the Gamma-ray Bursts Coordinates Network (GCN). By default, this initial version of the alert will not contain information about which observatories triggered, unless prior agreement has been reached otherwise. That information can be released in an updated version of the alert, after participants have had an opportunity to confer.

For each potential discovery, a team that includes representatives from each contributing observatory will confer to decide when and how to make the data public. If a unanimous decision is reached, the representatives may then request additional time to confer with colleagues in their respective collaborations. The length of this additional time period must be decided in advance and may be different for different alerts. The AMON team will keep track of this information in a location accessible to all observatories participating in the alert.

If unanimity remains, the alert will be approved and released by the participating observatories with help from the AMON development team. In phases 2 and 3 of AMON, this process will include revealing which observatories triggered. In the final phase of AMON, it will also include releasing the alert publicly via GCN. Where a unanimous decision can not be reached, data from the dissenting observatories will be removed from the coincidence analysis. If the subsequent alert is still of sufficient significance, it may be released without reference to the excluded data. Those observatories that did not participate in the alert may release their data independently, subject to the confidentiality conditions listed above.

Publication/Data Release

The data release procedure is illustrated in Fig. 1. A list of rules for data release is also provided below. These rules apply to all triggering and follow-up observatories that have signed the AMON MoU, for results produced via archival or real-time analysis.

1. At any time an observatory may independently release/publish its own data without discussing another observatory's data. If the AMON alert(s) have already been made public, or if the observatory can make a stand-alone claim for discovery (*i.e.*, if their signal can be statistically distinguished from background, independent of other observatories' data), the publication may reference the AMON alert ID(s). In the latter case, no other information from the AMON alert(s) other than the ID may be published, save that which is already in the public domain.

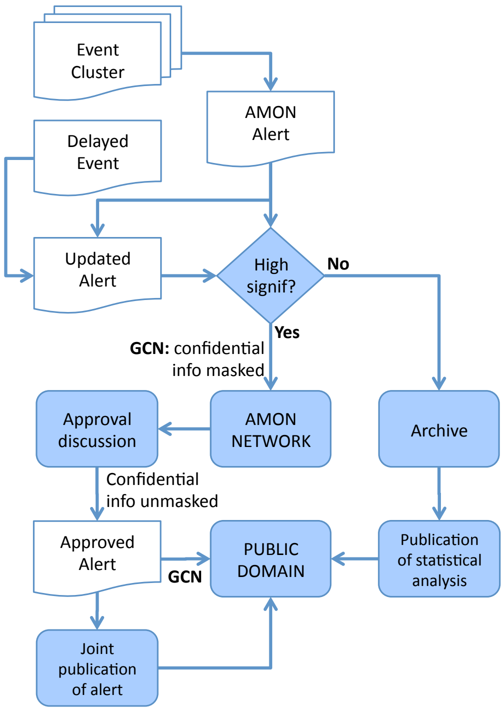

Figure 1: Data release during all phases of AMON will be carefully controlled. A cluster of events results in an AMON-generated alert which, if determined to be of high significance, is a candidate for immediate internal (non-public) distribution via GCN. Delayed events can subsequently be added to the analysis, enabling a review of whether an alert is of high enough significance for distribution. The initial in-network versions of the alert contain no reference to the triggering channels. However, high-significance alerts will initiate an immediate discussion within the AMON consortium, with the goal of releasing that information as soon as possible. Publications about the alert may follow at an appropriate time, as discussed in the text. Alerts that by themselves were not significant enough to be released through GCN are archived and may be studied by Consortium participants for stacked analyses and subsequent publications. These publications are expected to be statistical in nature and will not release individual alert information.

2. For joint AMON publications, unless decided differently by all participating collaborations, the author list will be strictly alphabetical by last name, with each name referenced by lower-case roman letter(s) to their collaboration(s) and by number(s) to their home institution(s). The ordering of the collaborations will follow that of the alphabetized author list. The ordering of the home institutions will be alphabetized independent of the author list. In addition, individuals of any collaboration who have contributed to analysis algorithms or AMON systems development may be included and referenced as members of the "AMON Development Team." *N.B.:* If all participating collaborations agree to use a different author list ordering, that ordering will be used instead of the ordering described above.

3. For joint AMON publications, the list of participating observatories will be established *a priori* by the configuration of the analysis algorithm (whether the analysis is real-time or archival). However, an observatory may be excluded from the author list if (a) they have already published a stand-alone discovery of the same astrophysical transient or (b) it is determined that they have no sensitivity to the temporal and spatial region of the study (*e.g.*, when a real-time alert does not occur within the observatory's field of view). In either case, individuals from the excluded observatory may still be included in the author list if they contributed to development of the associated analysis algorithm.

4. Whenever data from an observatory is referenced in a paper, the members of that observatory will be included in the author list unless the data is already in the public domain, even in the case where said observatory reports a non-detection. Specifically, this holds for follow-up observatories that may not have been included in the triggering analysis. In the case that an observatory's data is already in the public domain, individuals affiliated with that observatory who contribute to the published analysis will also be included in the author list and shown as affiliated with said observatory.

5. In the case of a joint publication based on real-time data, a best effort will be made to

include data from observatories that transmitted data to the AMON framework after some delay, whether the delay is due to systems latency, the time required to produce follow-up imaging, or the underlying astrophysics. However, the authors may proceed without the new data if inclusiveness would cause a significant delay in submission of the publication.

6. In the case where one or more observatories wish to delay release of their data, or if they have been unresponsive after a brief opt-in period, they may be removed from the author list of a joint publication and the analysis re-run without their data. If the repeated analysis is still statistically significant, the remaining authors may continue with the publication. Any observatory may subsequently release their own data independently.

References

[1] `http://amon.gravity.psu.edu/participants.shtml`.

[2] Sarah Emery Bunn, Roy Williams, Rob Seaman, *Hotwiring the Transient Universe*. 2010.